Visualization Support for Training PINNs as PDE Solvers

Afrah Farea[1], Yiwen Xing[1], Saiful Khan[2], and Min Chen[1]
[1]Department of Engineering Science, University of Oxford, 7 Keble Rd, Oxford OX1 3QG, UK
[2]Scientific Computing, Rutherford Appleton Laboratory, Science and Technology Facilities Council (STFC), OX11 0QX, UK

1 Introduction

Physics-Informed Neural Networks (PINNs) offer a mesh-free alternative for solving partial differential equations (PDEs) by embedding the governing equations directly into the calculation of training loss. Since the seminal work of Raissi et al [1], PINNs have been applied across many fields in science and engineering. However, the behaviours of PINNs are often difficult to understand or explain, such as where, why, and how they converge or do not, owing to the highly non-convex optimization of their composite loss, gradient imbalance, and uneven convergence across a spatial domain. Practitioners typically rely on global learning curves, which aggregate spatial information too quickly. In order to obverse temporal behaviour of a PINN in its spatial context, we use visualization to depict spatio-temporal behaviours. In particular, building on our information-theoretical understanding of visual multiplexing [2], we designed visualization plots to overlay multiple co-located scalar fields (e.g., local residuals, per-term loss gradients) within a single view. Such visualization reduces the pace of information loss caused by global learning curves, improving the interpretability and explainability of PINNs trained as PDE solvers as well as the effectiveness and efficiency in problem diagnosis and hyperparameter tuning.

2 Training PINNs as PDE Solvers

Training PINNs for solving PDEs involves many iterations (epochs) where the PINN under training attempts to estimate the solution from a set of sparsely sampled data points in the relevant spatio-temporal domain. The loss function in PINNs is typically defined as: $\mathcal{L}(\theta) = \arg\min_{\theta} \sum_{k=1}^{n} \lambda_k \mathcal{L}_k(\theta)$ where n denotes the total number of loss terms, $\mathcal{L}_k$ are the individual loss functions for each constraint with weighting coefficients λ_k, and θ represents the trainable parameters.

3 Visualization Support for Training PINNs

Traditionally, different variables are visualized in different heatmaps, which are difficult to compare [3, 4]. Because analysing multiple juxtaposed 2D heatmaps is cognitively demanding Fig.1(a), we utilize visual multiplexing [5], as shown in Fig.1(b), (c), and (d), which can depict multiple pieces of information associated with the same data points simultaneously. With superimposed visualization, computational scientists can intuitively compare fields without the need to shift their visual attention from one heatmap to another. In addition to heatmaps, we can also include contour plots and multivariate glyphs in multiplexed visualization.

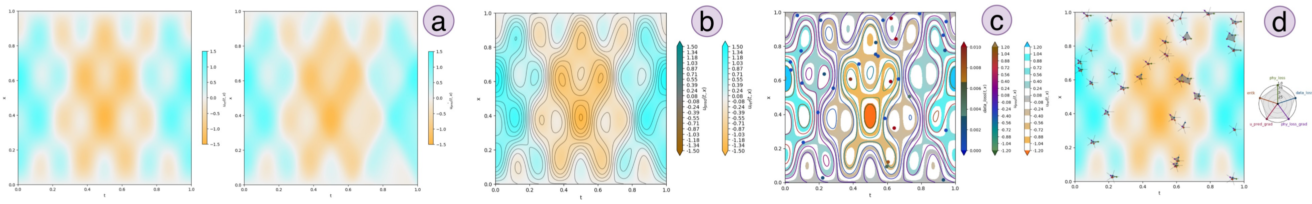

Figure 1: (a) Juxtaposed 2D scalar fields difficult to analyze, but with multiplexing we can visualize more variables, e.g., (b) two variables, (c) three variables, and (d) six variables. See [5] for details.

4 Conclusion

Providing effective visualization support for training PINNs allows computational scientists and machine learning researchers to observe more intermediate training data efficiently. It can also allow them to analyse, reason, and make decisions in development workflows for PINNs more effectively. We are conducting further research in automatic PDE discovery while developing a visualization-supported infrastructure for supporting a variety of machine learning workflows, including machine learning for PDEs.

This work is part of VIS4ML4HD project funded by UKRI/EPSRC grant EP/X029557/1.

References

[1] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

[2] M. Chen, S. Walton, K. Berger, J. Thiyagalingam, B. Duffy, H. Fang, C. Holloway, and A. E. Trefethen. Visual Multiplexing. *Computer Graphics Forum*, 33(3):241–250, 2014.

[3] Afrah Farea, Saiful Khan, and Mustafa Serdar Celebi. Multi-Objective Loss Balancing in Physics-Informed Neural Networks for Fluid Flow Applications. In *HiPC 2025*, pages 108–118, 2025.

[4] Afrah Farea and Mustafa Serdar Celebi. Learnable activation functions in physics-informed neural networks for solving partial differential equations. *Computer Physics Communications*, 315:109753, 2025.

[5] Yiwen Xing, Afrah Farea, Saiful Khan, and Min Chen. A Multiplexing Design Space: Theory, Method, and Application, 2026.